Expert systems: their role in Mineral Processing in the UK

P. Tucker* and K. A. Lewis*

Abstract
The need to promote and maintain high plant efficiencies is becoming vital for the UK minerals industry. Part of the solution must lie in the use of computer-based tools to facilitate better plant design, more optimal operation and more cost-effective maintenance. Although traditional computer-based tools (such as mass balancing, process modelling, simulation, on-line monitoring and conventional control) can go a long way to fulfilling these needs, they do not always provide the best solution. Expert systems offer significant potential in filling some of the gaps. The inter-relationship between expert system methodologies and traditional software techniques is discussed.

This paper examines the areas where expert systems could be most profitably applied within the UK minerals industry and reviews the (small number of) applications where they, and allied techniques, are already in commercial operation within the industry. The problems associated with system development, knowledge elicitation and industrial exploitation are discussed within a minerals industry framework and the likely benefits are detailed. The development of an expert system is illustrated with reference to two systems under development at Warren Spring Laboratory: The first is a diagnostic system for trouble-shooting shaking table operations and the second is a consultant system for advising on process routes for specific minerals. The diagnostic system is currently being evaluated within the Cornish tin industry. Experiences, in using the system, are reported.

Expert systems in the minerals industry

Expert systems are rapidly becoming an accepted part of modern technology. Certainly they are now receiving widespread attention in both technical and popular literature and most readers will be familiar, to some degree, with what expert systems represent and what they can hope to achieve. Rather than reproduce the arguments here, the reader is referred to references 1 to 3 where informative and detailed discussions can be found. It suffices here to provide a definition and partial glossary of some of the terminology involved. This paper is more concerned with the application of expert system techniques within the minerals processing industry. Firstly, however, the definition:

An expert system is a computer program that uses encoded knowledge and inference procedures to solve problems that normally require significant human expertise for their solution. An expert system will operate within a narrow area of expertise and will give the same type of advice as a human expert.

The term expert system is often used synonymously with the term 'knowledge-based system' though, more strictly, an expert system is a special kind of knowledge-based system that (i) uses knowledge acquired from an expert rather than from other sources

* Warren Spring Laboratory, Department of Trade and Industry

and (ii) can in some way explain its decisions to the user. Knowledge is captured and encoded ('knowledge engineering') from skilled people and experts in a specific field ('domain'). An expert system comprises the following elements: (i) a 'knowledge base' containing the facts, rules and heuristics specific to the domain, (ii) an 'inference engine' which can actively use the knowledge within the knowledge base and (iii) a user interface. A ready-built (integrated) inference engine and user interface is termed an 'expert-system shell'.

The early history of expert systems has close links with the minerals industry with the PROSPECTOR system for geological exploration,[4] well-known as one of the pioneers. Prospector made news in 1982 when it correctly predicted the existence of a molybdenum deposit (which geologists, given the same field data, had missed). Since then, many expert systems have been developed and applied both in exploration and mining areas, though relatively few are reported for the processing sector of the industry. Of these, many are aimed at process control applications. Melama *et al*[5] describe the implementation of an expert system to the control of the thickening section of the Outokumpu Kokkola zinc plant. The stated objectives of this project serve well to define the potential of expert systems in process management: (i) to be able to quickly obtain a rough overall perspective of the status of the process, (ii) to provide background supervision that prompts operational staff on important factors that may have been otherwise overlooked, (iii) to ensure that uniform procedures are implemented by all shifts when implementing control measures. Expert system technology is also in use in the Polaris Mine lead/zinc plant where a real-time controller has been developed for the flotation circuit.[6] Expert system control has also been developed and evaluated for the grinding circuit at the Silinjarvi concentrator.[7] In UK industry, perhaps the most notable expert system in regular use, is the LINKman supervisory control system developed by Blue Circle Industries PLC and Sira Ltd.[8, 9] The system, which is based on mimicking the manual actions of the plant operator, has been profitably installed to control five of Blue Circle's cement kilns and has also found application within the petroleum and glass industries. Quoted savings, for the production of 1.1 M tonnes of clinker, are $930,000/year. The payback period was 3 months. Also within the cement industry, Flamant *et al,*[10] report the development of a trouble shooting expert system for a clinker grinding mill. Other reported diagnostic expert system applications include on-stream XRF analysis,[5] dense media cyclones[11] and shaking tables[12] (and this paper). Flowsheeting systems are at an early stage of development. Marchal *et al*[13] report one for mica beneficiation. They also highlight the potential use of their system as a tutorial tool.

Practical experience at Warren Spring Laboratory

Expert system developments, at Warren Spring Laboratory, have been a logical extension to the laboratory's traditional role in providing services in flowsheet design and optimization of mineral treatment processes. In recent years, considerable effort has been devoted to the development of software tools to assist in these applications. Although the major developed techniques (modelling, simulation and material balancing) are now realising considerable benefit within the industry, [14, 15]) on their own they can not cover all applications. Simulation can assess flowsheets but it can not design them. It can optimize operations but is not well suited to fault diagnosis if the operation fails to meet expectations. Material balancing is invaluable in analysing how a flowsheet is operating, though it is a lengthy procedure unsuited to providing results in real time. Further it can not take direct account of qualitative information (eg, visual observations, which are often so important in formulating any judgement on plant performance).

Realising these deficiencies, Warren Spring Laboratory in conjunction with Carnon Consolidated Ltd and Beralt Tin and Wolfram S.A. (Portugal) set out to develop a more complete computer-based toolkit, in which expert systems would fill gaps left by the other two methods. The methodology adopted is shown schematically in Figure 1.

The one thing that the mineral processing industry is not short of is heuristic knowledge. This, in itself, makes it an ideal candidate for expert systems. Quite often this knowledge is expressed in the form of IF...THEN...ELSE rules (often termed 'Production Rules' – a standard form of knowledge representation in expert systems). Available rules range from hard fact supported by scientific theory, through empirical rules to widely (or locally) held rules of thumb. One of the most debated points in the whole field of expert systems is 'How then does one get these rules out of the expert?' There is, unfortunately, no universal answer to this question; the solution being more to do with psychology than technology. The way that Warren Spring Laboratory is tackling the problem is as follows:

The knowledge engineers at Warren Spring Laboratory, by the very nature of their past experience, are themselves minor experts in the domain. Their basic knowledge is readily supplemented through published literature and historical reports held at the laboratory. By utilizing this data, a fairly comprehensive prototype system can be built. During this initial prototyping phase, clarification (as necessary) is obtained in-house through consultation with personnel who are more expert within the specific domain. These same experts are then used to validate the prototype system and, in the process, are able to contribute further knowledge towards the system. A second-level prototype

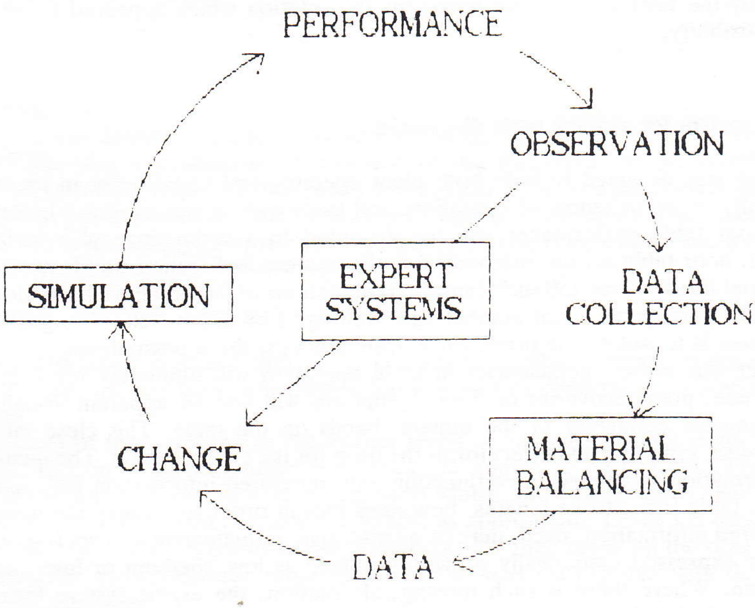

FIGURE 1
A computer-based methodology for an operating plant

is then constructed and distributed to the potential industrial user. Again the system is validated with the industrial expert contributing his deeper and/or local knowledge into the system. This approach seems to offer distinct advantages: The prototyping is rapid with a major part of the groundwork knowledge compilation being undertaken by the software developers themselves (who know exactly what knowledge is required). The experts can then be given a model set of rules which they can evaluate (or criticise) and also use as a template for contributing their own deeper knowledge . . . at their own convenience. One criticism which could be levelled at this approach is that too much tacit knowledge might be taken for granted. However, it must be borne in mind that the end user will also hold much the same tacit knowledge (eg, the plant operator should know full well what 'increase the pulp density' means without needing to have it explained to him). The systems being developed are not aimed at a lay user. (A graphical description on the importance of tacit knowledge is given by Collins.[16])

A second debating point in expert systems is what computer language or expert system shell to use for the development. The choice here is rarely crucial though, if the choice is reasonable, it will simplify the task involved. It has already been stated that the domain knowledge is well represented by a production-rule formalization. Further, the problems tackled to date seem well suited to 'Goal-driven reasoning', that is testing a conclusion by seeing if the facts and rules support that conclusion. These features are particularly suited to a PROLOG implementation, and that language has been chosen for all developments so far. Proprietry shells had been considered at the outset of the work. Although offering an easier development environment, they tend to lack the flexibility and generally impose more limitations than do pure languages. At that time, because we did not then know the eventual sizes or demands that would be made by the final systems, we opted for the solution which appeared to offer the greatest flexibility.

An expert system for shaking table diagnostics

Overview

This system was designed to help both plant operator and metallurgist in identifying tabling faults, or inefficiences of operation, and to provide advice on possible remedial actions. Poor table performance can be attributed to a wide range of contributory factors (eg, poor table set-up, mechanical faults or wear and tear or problems with the feed material itself). Over 150 such causes (combinations of factors) have been identified to date, with perhaps an equal number still 'waiting to be discovered'. The goal of the expert system is to isolate the prime cause from amongst these possibilities.

Whatever this cause, inefficiencies in table operation will ultimately result in poor product grade, poor recoveries or both. Symptoms will also be apparent visually, for example unusual behaviour of the mineral bands on the table. This close interplay between cause, symptom and effect forms the basis for the expert system. The qualitative visual information is applied in conjunction with measured information (eg, achieved recoveries, table tilts, stroke lengths, flow rates etc) in order to achieve the diagnosis. The measured information, itself, may be quantitative or qualitative in form (eg, strokes per minute expressed numerically or simply judged as low, medium or high) or may be unknown. Where there is such missing information, the expert system forms the 'best' diagnosis on the information that it has got. If, however, the information is particularly sparse, then the system may only be able to narrow down the list of possibilities to the two or three most likely causes.

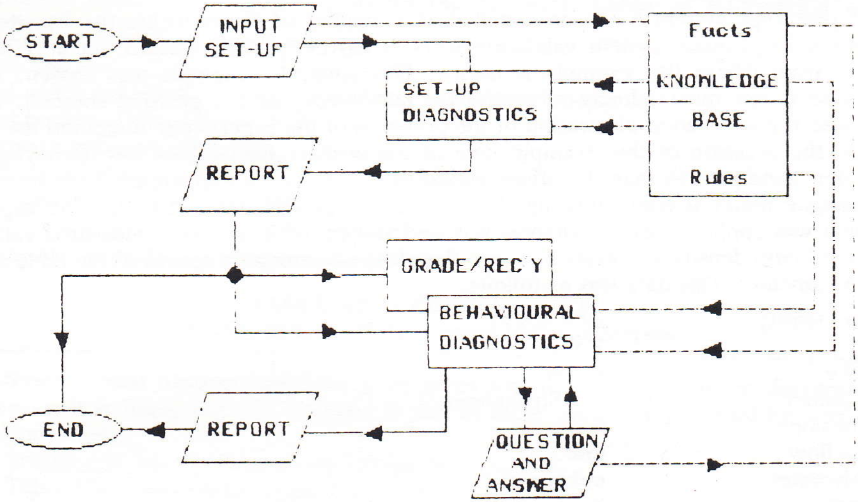

FIGURE 2
System structure for a diagnostic expert system

The solid lines show the path through the system. The dotted lines represent the data flow.

System Structure

The system was designed around three levels of diagnostics (Figure 2). The entry level (level 1) provides a preliminary assessment of the measured, or perceived, set up parameters (strokes, frequency, tilts, wash-water) against the feed parameters (flowrates, pulp density, particle size) and the table duty (roughing or cleaning). The expert system rules, at this level, simply comprise the 'normal' operating ranges and some well-known heuristics, eg, roughing operations require:

(more water, more ore, less tilt, longer stroke)

A diagnosis report is issued after the level 1 analysis, either warning on possible fault conditions (eg, tilt too high, wash-water too low etc) or confirming no obvious faults in the set-up.

Diagnosis can then proceed to the deeper levels of analysis. Level 2 starts from the basis that measured grades and/or recoveries are known to be too low. The system then prompts for the visual indicators (eg, band widths, drifting bands etc) and requests other information (eg, feed preparation details) as appropriate. Level 3 is chosen when the grade/recovery data is not available. Diagnoses are then made on the observational data only. The final diagnostic report provides the fault list together with the reasons behind the diagnosis.

A full description of the logical reasoning, within the system, is outside the scope of this paper. Full details can, however, be found in reference 11. This paper instead concentrates on the application of the system in the UK minerals industry.

A practical example

The prototype system has been installed at a number of plants within the European Community. Formal system validation work is currently being carried out at Wheal Jane, from where this example is drawn. This particular example was chosen, not because it was outstandingly successful (in fact it only gave a guarded success), but because it gives a good illustration of the potential of the expert system method itself.

On the occasion of this example, one of the primary fines tables was displaying a 'fuzzier' band pattern than the others within the bank, with an apparently high amount of gangue material contaminating the concentrate product (see Plate I). The expert system was applied. Level 1 analysis was undertaken using, as input, measured values for feed pulp density and mass flow and the plant metallurgist's record of the table set-up parameters. This data was as follows:

Pulp Density	:	29% sol/wt
Tilt	:	1.5°
Stroke	:	8 mm
Frequency	:	274 Stroke/min
Long-slope	:	0.3°
Mass flow	:	low
Wash-water	:	unknown
Particle size	:	fine
Duty	:	roughing

Level 1 diagnostics reported 'No errors in table set-up detected so far'.

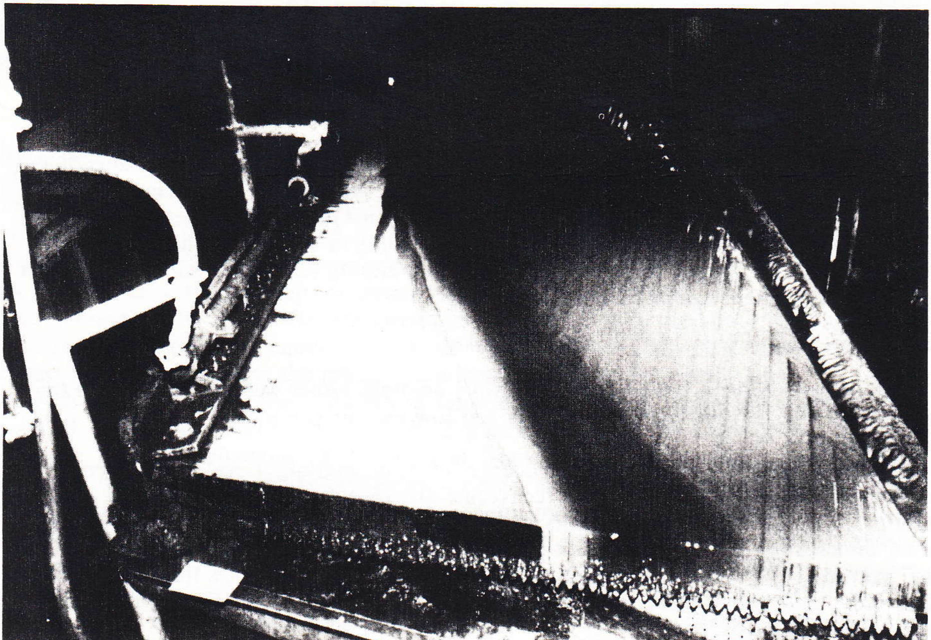

PLATE I
Shaking table performance 'as found'

Analysis proceeded to Level 2 where the following additional information was supplied:

Grade too low	
Band width unknown	(ie, bands fuzzy and not well defined)
Band position OK	(*)
Recovery unknown	
Feed classified	(by cyclones and hydrosizer)
Feed deslimed	(by cyclone)

The level 2 analysis report is reproduced below:

TABLE AID – FINAL REPORT

The final analysis of table operation is as follows:

Insufficient visual observations have been given for the table fault to be diagnosed. However, it is likely that the problem is due to table overloading caused by lengthy residence times.

The possibilities have been narrowed down to:

Table TILT may be too high.

FLUID FLOW may be too low: increase wash water or decrease pulp density.

Table FREQUENCY may be too low.

Data from level 1 indicates that the frequency is probably OK

Press R for reason why, C for hard copy, Q to quit

It is seen that a short list of three possible faults was identified. The first was immediately discounted as highly unlikely. Whilst not in direct contradiction with the system rules, it appeared implausible and perhaps, more than anything, pointed to an insufficiency of rules within the system itself. The third possibility appeared even stranger, being the result of two conflicting diagnoses. Perhaps, therefore, a low fluid flow was the real cause underlying the observed poor performance. In fact, increasing the fluid flow (by increasing wash-water) only gave a marginal improvement and subsequent analysis returned exactly the same diagnosis. The major fault was, in fact, identified as slippage in the drive belt. This had the effect of reducing the drive frequency below its perceived setting! The true diagnosis is in fact given by the third reported possibility. The apparent conflict, which is now logically explained, was in fact the key. In conclusion, the diagnosis was able to pinpoint the location of the fault (if not its root cause) and further identified low table water as a contributory factor. On repairing the identified fault, the observed table performance improved markedly (Plate II) where it is seen that the wider, sharper concentrate bands led to a much higher cassiterite grade being recovered.

The example also furnished additional information, and therefore expert system rules, and allowed a new more definite diagnosis to be added to the system. This aspect and the ready way it can be implemented is an integral feature of any expert system. The system can grow via continuous refinement, thus realizing an increased potential the more it is used.

(*) This was the 'best' judgement of the authors, at the time of validation. Subsequent assessment by Wheal Jane staff was that the band position was 'too low'.

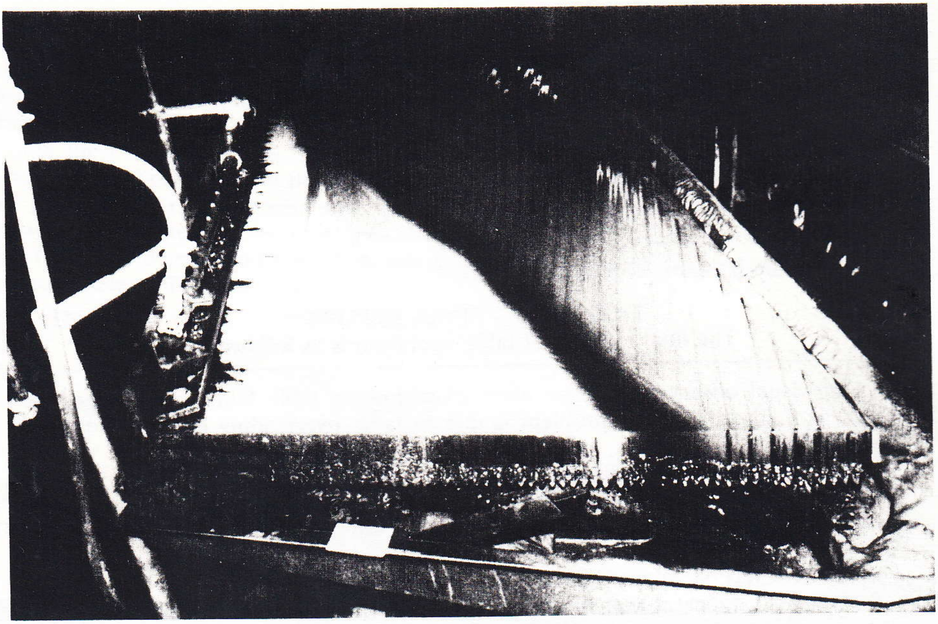

PLATE II
Shaking table performance after fault correction

An expert system for gold processing
This system is designed to complement the traditional role of Warren Spring Laboratory in the development of process flowsheets for specific minerals. The system integrates the in-house expertise in both mineral processing and extractive metallurgy. The system is still very much in the prototype stage, both in its current level of development and in the structural concepts that are being developed. The system is viewed as a proving ground to ascertain feasibility and to underpin any future development of expert flowsheeting systems.

The system is being designed around two levels of analysis. Firstly, the best 'type' of flowsheet is identified. At this level, the flowsheet is simply a pre-defined network of stages (eg, comminution followed by gravity concentration and cyanide leach could be a simple three-stage flowsheet). The second level of analysis serves to define the best (pre-defined) network of (one or more) devices, or processes, for each of the processing stages within the overall flowsheet (see Figure 3). For example, one processing stage is 'Removal of carbonaceous material'. The stage options here include 'Oxidation', 'Rendering it inert', 'Roasting' or 'flotation'. Similarly the options 'Logwash and deslime', 'Trommel and deslime' or 'No action' could be selected for a 'Clay removal' stage, depending on whether the ore is sticky, not sticky or if clays are present at all.

Each stage is represented as a series of attributes, (eg, product specification, product value, resource needs etc) and a list of the component devices or sub-processes. Similarly each included sub-process may be assigned its own set of attributes. The function of the expert system is to assess these attributes against the local (user-defined) constraints and to select the optimum flowsheet, and stages therein, to best meet the imposed constraints. The system will eventually have the potential to rank alternatives and to justify the reasoning behind its selection.

The future

The concepts being developed on the gold expert system open up a number of exciting opportunities. The first is a formal integration of the expert system methods with relational database technology. It is not difficult, conceptually, to see how a major part of the knowledge base (formulated in terms of stage and device attributes) could be contained in, and accessed from, the structure storage facilities offered by a database. In this way, new classes of attributes (eg, environmental and safety) could be readily incorporated and used in the assessments. A second possibility lies in the integration of conventional simulation and modelling techniques with the expert system. Here, it could be envisaged that the modularity encompassed by the system (that is its representation as a network of stages) offers an additional potential for treatment in a more quantitative manner. These two ideas are embodied in a new proposal, by Warren Spring Laboratory, for a multi-company collaborative project, under the acronym IMPREST (Industrial Mineral Processing Rule-base for Environmentally Safe Technology). The aims of IMPREST are summarized in Figure 4.

Several other expert system initiatives are currently being considered by Warren Spring Laboratory. One worthy of note (in the context of the UK minerals industry) is the proposed development, in conjunction with Carnon Consolidated Ltd, of an expert

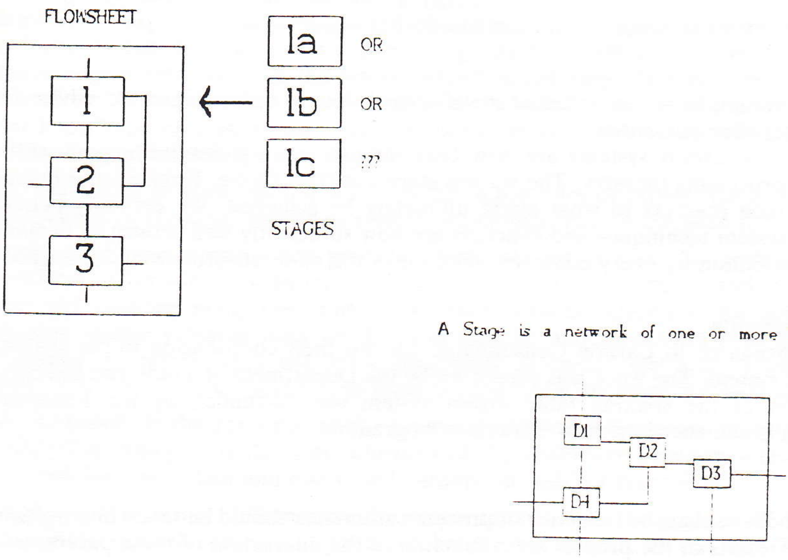

FIGURE 3
Concepts behind a flowsheeting expert system

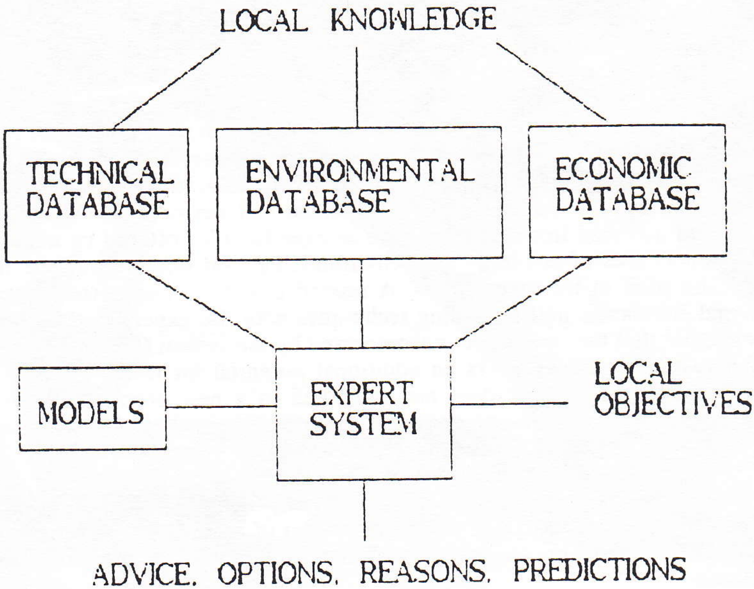

FIGURE 4
Schematic of a full flowsheeting expert system

supervisory system to assess collected metallurgical data in order to provide advice for the plant operating personnel.

In conclusion, expert systems are now beginning to make a positive impact within the mineral processing industry. The success stories at Blue Circle, Kokkola and Polaris Mine give some pointers to what might ultimately be achieved. We certainly believe that expert system techniques and concepts are now sufficiently well proven to demand the closest attention by every company involved in the minerals processing industries.

Acknowledgements
Particular thanks go to Carnon Consolidated Ltd for their contribution to the shaking table expert system. The work was supported by the Department of Trade and Industry. Development of the shaking table expert system was co-funded by the European Commission under the third Raw Materials Programme.

Discussion

D. Daka: I believe that the influence of process parameters should be taken into account because the results of the process are a function of the interaction of these parameters. Have any expert systems for flotation been implemented?

Authors: The only example, of which we are aware, where an expert system for flotation has been implemented on a process plant is at Polaris Mine, the reference for which is given in the text of the paper.

The flotation process is perhaps the prime example of where numerical prediction methods (eg, modelling) have so far failed to provide a general solution. In part, the reason is that much of the knowledge needed to describe the process cannot be expressed as numerical algorithms. Expert systems would seem to offer the potential to resolve this problem. Whether or not heuristics can be identified to account for the influence of process parameters, in the general case, is however not yet answered.

R. Klymowsky: In the discussion following your paper you mentioned that expert systems for flowsheet design of gold processes will be available in the near term. Could you be more specific about what you mean by 'near term'?

Authors: The gold processing system has been built purely as a prototype in order to establish and validate the concepts and methods needed for a workable flowsheeting expert system. As such, the system is not in a commercial form though it could be made commercially available within the next year provided sufficient industrial interest is shown.

J Zhang: Expert systems are very useful for educational purposes but the difficulties in programming human knowledge of mineral processing make them less acceptable for operating purposes. Modelling and simulation are far better techniques for solving problems in industry.

Authors: We agree that modelling and simulation are extremely valuable in helping to solve problems in mineral processing, as the paper 'Computer simulation at Wheal Jane', presented at this conference, clearly demonstrates. Valuable as such methods are, however, they cannot solve all problems, particularly if the solution requires heuristic knowledge. We are strong advocates of using all possible tools in order to achieve the best solution and we see an integrated approach involving modelling, simulation, expert systems and indeed other computer methods as providing a comprehensive toolkit for the mineral processing engineer. Of course, as you point out, knowledge elicitation from the human expert is not easy. This is a common problem for all expert systems and is not peculiar to the mineral processing industry. Methods for knowledge elicitation can, and have been, devised and successfully applied over a wide range of knowledge domains.

H. Lofthouse: In your opinion what are the future roles of expert systems? Are they to be used as operational aids or as systems which will remove the need for operators?

Authors: Expert systems will provide a valuable set of tools for both engineer and operator. They will assist by (i) enabling the human to make better and faster decisions and (ii) enabling many decisions to be made without referral to the specialist, thus freeing his/her valuable time which could then be allocated more profitably. Like all software, therefore, they will serve principally as operational aids. It would be very dangerous, for obvious reasons, to replace completely operators by such systems.

P. A. Dowd: Is the response of the system qualified by the certainty/reliability of the data? For example, is the user informed of the difference between a response based entirely on 'hard' data and one based entirely on 'soft' or fuzzy data?

Authors: We do not differentiate the response solely on the grounds of 'hard' versus 'fuzzy' data. Each response is generally based on a mixture of the two. The certainty of the response is determined by the amount of data in support of the conclusion and whether any of the data are conflicting. Conflicts, where hard data appear to contradict qualitative observations, are highlighted. In many cases such conflicts arise because the hard, rather than the soft, data are in error, either through measurement error or as a result of undetected equipment malfunction. These are precisely some of the faults that the system is designed to trap.

References

[1] Hayes-Roth, F., The knowledge-based expert system: A tutorial. IEEE Computer (Sept 1984) pp 11–28.
[2] Gervarter, W. B., Introduction to artificial intelligence. Chemical Engineering Progress, 83 (Sept 1987) pp 20–37.
[3] Waterman, D. A., A guide to expert systems. Addison-Wesley: Reading, Mass. (1986).
[4] Duda, R., Gaschnig J. and Hart, P., Model design in PROSPECTOR consultant system for mineral exploration. ESMA (1979) p 153.
[5] Melama, H., Barker, D., Hirvonen, M. and Hartikainen, J., Experiences gained from using an expert system approach on process management and control applications., Proc Artificial intelligence in minerals and materials technology conference. Univ Alabama (Oct 1987), pp 66–94.
[6] Harris, C., Kosick, G., Expert system technology at the Polaris Mine. Proc. Annual conference of the Canadian Mineral Processors, Ottawa (Jan 1988).
[7] Jamsa, S. L., Experiences in the use of expert system in grinding process control in Silinjarvi concentrator. Proc XVI IMPC, Stockholm (June 1988). Elsevier: Amsterdam, 1988, pp 1759–1770.
[8] Haspel, D. W., Lorimer A. D. J., Southan C. J., Taylor R. A., Blue Circle high level kiln control., Proc IEEE conference, San Francisco (May 1987).
[9] Haspel D. W., Southan, C. J., Taylor R. A., The benefits of kiln optimization using LINKman and high level control strategies. World Cement, 18 No 6 (July 1987).
[10] Vanderstichelen, J., Expert system prototype for process diagnosis in the clinker grinding circuit of a cement factory. Proc CANMET symposium on expert systems in the mineral industry, Ottawa (March 1987).
[11] Layton, R., Hunt, J., Reeves, T., Dense medium separation diagnosis expert system: Prototype evolution and production model design. (to be presented at APCOM 1989).
[12] Tucker, P., Lewis, K. A., An expert system for shaking table diagnostics. Mineral Engineering, 2, No 1 (in press).
[13] Marchal, J. H., Kleikamp, D. J., Sullivan, G. V., Wilkerson, R. W., An expert system for mica beneficiation. Proc Artificial intelligence in minerals and materials technology conference. Univ Alabama (Oct 1987), pp 142–148.
[14] Lewis, K. A., Wells, A., Tucker, P., Computer simulation at Wheal Jane, Proc Mineral processing in the United Kingdom, Leeds 1989: (this volume).
[15] Newton, S., Computer mass balancing for plant process audit, Proc Mineral processing in the United Kingdom, Leeds 1989: (this volume).
[16] Collins, H. M., Domains in which expert systems could succeed. Proc third international expert systems conference (June 1987), pp 201–206. Learned Information: Oxford, 1987.